

REMARKS

The Patent Office is again respectfully requested to change its records to reflect the new Power of Attorney that was filed on March 20, 2008 and acknowledged in the Office Action on September 30, 2009. Per the Power of Attorney, the Patent Office should be corresponding with

Customer No. 021005

David J. Thibodeau, Jr.

Hamilton, Brook, Smith & Reynolds, P.C.

530 Virginia Road

P.O. Box 9133

Concord, MA 01742-9133

Claims 1-2, 8-10 and 22-27 stand rejected under 35 U.S.C. 103(a) as being obvious over Dimitri (U.S. Patent 6,839,802) in view of Jacobsen (U.S. Patent 5,392,244) and Kim (U.S. Patent 6,717,436).

In aspects of the present invention, a performance process or system measures the actual performance of different locations in a storage array by making experimental read and write operations across a logical block name space. These measurements are used to determine whether various locations can be aggregated into like regions. (See for example, Applicants' specification at page 7, lines 4-11). A mapping process or element then aggregates the logical block names for locations measured to have an identical level of performance. This results in the creation of different storage pools. The Applicants' claimed process and system then assigns different RAID levels to these different regions based upon the different measured levels of performance of the locations within the regions (see at least the Applicants' specification at page 2, lines 17 through 19, and page 4 lines 9 through 11, and Fig. 4 and claim 1, 11 and 19).

Two different clients thus accessing such a system can utilize this storage pool and select different classes of both RAID level and different performance from the very same storage array (see the specification at page 11, lines 5 through 8), with both clients storing their data on the very same set of physical disks (see the specification at page 9, line 22 through page 10 line 2).

These storage array features and the attendant advantages are neither taught nor suggested by Dimitri, Jacobsen, and Kim references even when considered in combination.

In Dimitri, a storage medium is formatted into a plurality of “zones”. Each zone comprises one or more circumferential tracks, with each track comprising a plurality of addressable sectors. The innermost zones have fewer sectors than the outermost zones. When a request is received to write a file to the storage medium, a “utilization factor” is determined for the file and a zone for the file is selected. The determined file utilization factor is an expected rate at which the file will need to be accessed by the application, not the measured performance of the medium itself. The file is then written to the selected zone. (See Dimitri, the Abstract). Thus, files with relatively low access rates are assigned to low performance areas, and files with relatively high access rates are assigned to relatively high performance areas on the storage medium.

Applicants must therefore disagree with the Examiner’s characterization of Fig. 4 of Dimitri. That figure is not “logic for determining a level of performance for storage locations.” It is merely a table that describes the as-manufactured physical structure for the storage regions on the disk, including a number of sectors per revolution (revolutions per minute which are indicated as being the same for each zone) and a read rate in megabytes per second for the zone. It does not suggest performing a process as claimed for determining a measured level of performance for a plurality of storage locations. Nor does it suggest the claimed determining a level of performance by experimental read and write operations across a logical block main space, and then using those measurements to aggregate various locations by a logical block name.

Furthermore, the partitioning suggested by Fig. 4 in Dimitri is merely a predetermined partition based on the physical layout of storage areas on the disk, and not determined by measured levels of performance.

For these reasons alone, the rejection under 35 U.S.C. 103(a) is prima facie deficient and should be withdrawn.

Furthermore, at Column 8, lines 31-43, Dimitri is merely suggesting that a RAID controller would then write all data stripes to the same zone on different disks. This is because Dimitri is trying to ensure that the data transfer rate for all of the disks is the same, thereby

avoiding a situation where the RAID controller performance is limited to the performance of the inner most zone of all the disks to which data is striped. But, this does not amount to Applicant's claimed mapping process to map partitioned regions of the storage locations and aggregating logical block names having an identical level of performance to a selected section of the logical block named space. Dimitri mentions nothing whatsoever about logical block name spaces, and cannot therefore also suggest anything about aggregating logical block name spaces.

Jacobson does disclose a RAID controller. However, item 26 in Jacobson is just a first set of disks that stores original data; duplicative redundant data is then stored on a paired second set of disks 28. See Column 3, lines 40-42 of Jacobson. Figure 1 of Jacobson just shows a group of disks arranged in a mirror group 18 of multiple disks 20 and a parity group 22 of multiple disks 24. This type of redundancy is considered to be a RAID level 1 or "disk mirroring" type implementation. Thus, the Examiner's reading of Jacobson appears to be less than accurate since there is no suggestion in Jacobson that different RAID levels be serviced by the same physical disk.

Finally, Kim does describe a way to manage dynamic resizing of a logical volume. Kim also describes various RAID levels and that disk striping can be provided such that a single record can span multiple disks. Kim also mentions that better performance requires establishing a stripe that is wide enough to hold a typical maximum sized record. (See Column 1, lines 55-65 in Kim). In Figure 2, Kim does show a logical volume that may have disks 1, 2, 3, and 4 divided into four partitions and disks 5, 6, 7, and 8 each having one partition. Column 7, lines 7-15 of Kim also does suggest that after a disk partition is created using an operating system tool and a logical RAID volume can be constructed. However, there is no suggestion Kim of pairing this RAID assignment to measured performance levels.

In further aspects of Kim (at Column 8, lines 32-44), a logical volume map 72 includes information for identifying a logical volume, its extent size, a total number of extents in a corresponding volume, a stripe size, a RAID level, and a number of disk partitions. These are used to then construct the corresponding volume. However, this also does not amount to a suggestion of aggregating logical block names with the same measured performance criteria.

Therefore, even if one incorporated the storage controller of Jacobson and Kim into Dimitri's system, one would not arrive at the claimed invention. One would not end up with a

process that (a) measures performance of storage locations, (b) maps regions of like measured performance together as partitioned regions, and (c) aggregates their logical block names to provide different RAID level configurations on the same storage device.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By /David J. Thibodeau, Jr., Reg. No. 31671/

David J. Thibodeau, Jr.

Registration No. 31,671

Telephone: (978) 341-0036

Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date: September 17, 2010